Contribution ID: **66** Type: **Poster**

# Development of noSQL data storage for the ATLAS PanDA Monitoring System

For several years the PanDA Workload Management System has been the basis for distributed production and analysis for the ATLAS experiment at the LHC. Since the start of data taking PanDA usage has ramped up steadily, typically exceeding 500k completed jobs per day by June 2011. The associated monitoring data volume has been rising as well, to levels that present a new set of challenges in the areas of database scalability and monitoring system performance and efficiency. These challenges are being met with a R&D effort aimed at implementing a scalable and efficient monitoring data storage based on a noSQL solution (Cassandra). We present our motivations for using this technology, as well as data design and the techniques used for efficient indexing of the data. We also discuss the hardware requirements as they were determined by testing with actual data and realistic loads.

**Primary author:** Dr POTEKHIN, Maxim (Brookhaven National Laboratory)

**Presenter:** Dr POTEKHIN, Maxim (Brookhaven National Laboratory)

**Track Classification:** Track 1: Computing Technology for Physics Research